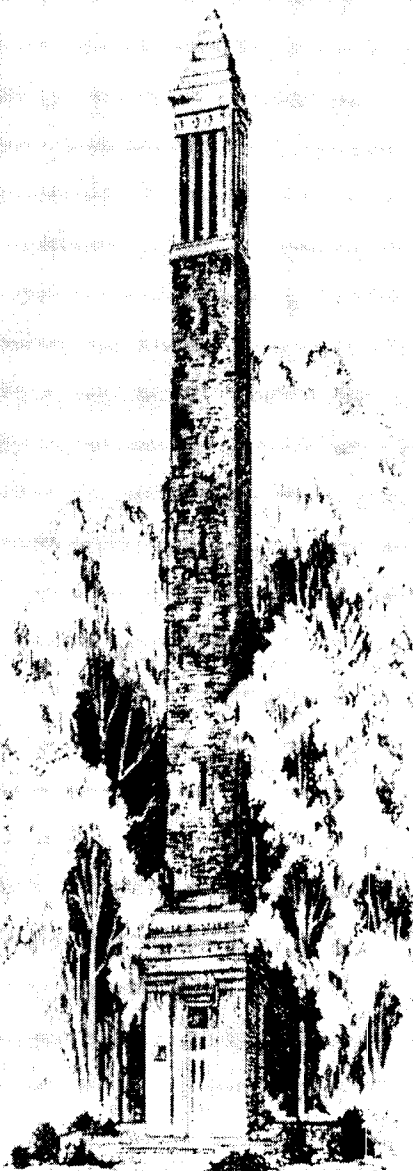


THE UNIVERSITY OF
ALABAMA

FOUNDED 1831



Final Project Report

*An Immunized Aircraft Maneuver
Selection System*

*Completed for NASA Ames Research
Center*

NMC 2-1326

Department of Aerospace Engineering
and Mechanics

The University of Alabama

205 Hardaway Hall

Box 870280

(205) 348-0066 (voice)

(205) 348-7240 (fax)

ckarr@coe.eng.ua.edu

July 2003

CASI

FINAL PROJECT REPORT

An Immunized Aircraft Maneuver Selection System

PREPARED FOR

NASA Ames Research Center
Moffett Field, CA 94035-1000

ATTENTION: Dr. Kalmanje KrishnaKumar

PREPARED BY

Charles L. Karr
Aerospace Engineering and Mechanics Department
The University of Alabama
Box 870280
Tuscaloosa, AL 35487-0280

July 2003

An Immunized Aircraft Maneuver Selection System

TABLE OF CONTENTS

1.0 ABSTRACT	3
2.0 INTRODUCTION	3
3.0 YEAR ONE STATUS	5
3.1 Task 1: Determination of the Building Block Architectures	7
3.2 Task 2: Identification of Immune System Features	8
3.3 Task 3: Choosing Suitable Test Functions	8
3.4 Task 4: Software Development for a Generic Immune System Based Optimizer	8
3.5 Task 5: Integration and Testing Using Vehicle Models	8
3.6 Summary of Status of Research Tasks	9
4.0 BASICS OF THE NATURAL IMMUNE SYSTEM	9
4.1 Humoral Immunity and Cell-Mediated Immunity	10
4.2 Cells of the Immune System	11
4.3 Clonal Selection is Basis for Specificity and Diversity	12
4.4 Memory Cells for Secondary Immune Responses	12
4.5 Molecular Markers Provide Self/Non-self Recognition	13
4.6 B Cells Generate Specific Antibodies	14
4.7 The Activation of B Cells	14
4.8 T-Dependent and T-Independent Antigens	15
4.9 The Molecular Basis of Antigen-Antibody Specificity	15
4.10 How Antibodies Work	17
4.11 The Activation of T Cells	18
4.12 How Cytotoxic T Cells Work	19
4.13 The Immune System's Capacity to Distinguish Self from Non-Self	20
5.0 ARTIFICIAL IMMUNE SYSTEM	20
5.1 Key Features of the Natural Immune System	20
5.1.1 Specificity	20
5.1.2 Diversity	21
5.1.3 Memory	21
5.1.4 Self/Non-Self Recognition	21
5.2 Key Features of an Artificial Immune System	21
5.3 Architecture for an Artificial Immune System	22

6.0 SYSTEM IDENTIFICATION PROBLEM	24
6.1 Overview of the System Identification Problem	25
6.2 Self vs. Non-self in the System Identification Problem	26
6.3 Receptors for Monitoring System Performance	28
6.4 Receptor-Antibody Pairs	29
6.5 Evolutionary Algorithm for Performing Clonal Selection	30
6.6 Results of Using the AIS on a System Identification Problem	31
7.0 STEPS FOR COMPLETING THE PROJECT	35
8.0 REFERENCES	35

An Immunized Aircraft Maneuver Selection System

1.0 ABSTRACT

The objective of this project, as stated in the original proposal, was “to develop an immunized aircraft maneuver selection (IAMS) system.” The IAMS system was to be composed of computational and informational building blocks that resemble structures in natural immune systems. The ultimate goal of the project was to develop a software package that could be flight tested on aircraft models. This report describes the work performed in the first year of what was to have been a two year project. This report also describes efforts that would have been made in the final year to have completed the project, had it been continued for the final year. After introductory material is provided in Section 2, the end-of-year-one status of the effort is discussed in Section 3. The remainder of the report provides an accounting of first year efforts. Section 4 provides background information on natural immune systems while Section 5 describes a generic architecture developed for use in the IAMS. Section 6 describes the application of the architecture to a system identification problem. Finally, Section 7 describes steps necessary for completing the project.

2.0 INTRODUCTION

Biologically inspired systems are becoming increasingly popular as methods for solving complex engineering problems. Among these systems are *artificial neural networks* (Fiesler and Beale, 1997), *evolutionary algorithms* (Bäck, Fogel, and Michalewicz, 1997), *DNA computation* (Jungheui and Deaton, 2001), and now *artificial immune systems* (Hofmeyr and Forrest, 1999; Hofmeyr, 2000). The immune system is a complex set of molecules, cells, and organs capable of performing numerous tasks including pattern recognition, noise tolerance, generalization, and optimization. The structures and underlying principles of natural immune systems are currently being used to develop new computational paradigms. These paradigms utilize information extracted from the problems themselves to form solutions. Although *artificial immune systems* have experienced some success (Dasgupta, 1999), no general formal framework for the approach has been developed.

Artificial immune systems have arguably been successful because they have a number of desirable properties, including:

- uniqueness: the immune system of each individual is unique and therefore both capabilities and vulnerabilities differ from one system to the next;
- distributed detection: the detectors used by the immune system are small and efficient, are highly distributed, and are not subject to centralized control or coordination;
- imperfect detection: by not requiring absolute detection of every pathogen, the immune system is more flexible – the body can trade off resources used on protection for comprehensiveness of coverage;
- anomaly detection: the immune system can detect and react to pathogens that the body has never before encountered;

- learning and memory (adaptability): the immune system can learn the structures of pathogens, and remember those structures, so that future responses to the pathogens can be much faster.

These properties result in a system that is scalable, resilient to subversion, robust, very flexible, and that degrades gracefully. The characteristics cited above are keys to natural immune systems' ability to solve complex problems. These properties contribute to four critical attributes possessed by virtually all artificial immune systems: (1) specificity, (2) diversity, (3) memory, and (4) self/non-self recognition.

Natural immune systems are highly complex and fascinating computational systems that: solve a unique problem in a specialized way, are highly distributed, have no centralized control, use learning and memory, and, are tolerant of themselves. Intelligent computer systems are quickly becoming so complex that we are fast losing the ability to analyze and understand them in any detail. Developing solutions to solve the problems that we are encountering is becoming an increasingly difficult task. As the complexity of our systems approaches that of biological systems, it is only natural that we turn to biology for inspiration, for biological systems can provide a vast amount of knowledge developed over eons of evolutionary experimentation. The immune system is a particularly sophisticated biological system, and one that has the potential to provide inspiration for solving complex machine learning problems.

Automated flight controllers are prime examples of complex systems that potentially require the capabilities afforded by artificial immune systems (AISs). An effective autopilot must be able ultimately to perform many of the tasks currently accomplished by pilots and operators. To develop such a system will require the investigation of numerous core capabilities including internal management of a vehicle's health, assessment of the current operating environment, and the determination of action in the form of strategic planning and tactical maneuver selection. It is the tactical maneuver selection that forms the crux of the current effort.

An AIS approach to the development of an intelligent control system will involve defining and learning computational or informational building blocks. As proposed for this project, these building blocks were to have been used to select optimal flight maneuvers, and were to have been developed both off-line and on-line. The building blocks developed off-line were to have consisted of a priori knowledge about flight maneuver selection. The building blocks developed on-line were to have been developed through the use of evolutionary algorithms, and were ultimately to have allowed for the selection of highly-effective flight maneuvers. In addition, the IAMS system was to have possessed the capability of employing numerous representational paradigms such as decision trees, neural networks, fuzzy systems, and look-up tables. The development of generic computer software for implementing the IAMS system described later in this report will provide a clear indication of the potential effectiveness of employing the artificial immune system approach, and has the potential to serve as a framework for extremely robust intelligent control systems.

This report describes progress made on the effort in the first year of what was to have been a two year project to develop an IAMS. It provides background information on natural immune systems, describes a general-purpose architecture for implementing an AIS, presents some preliminary results in which the AIS has been applied to a simple system identification problem, and outlines the steps necessary to successfully complete the project should it be continued at a later date.

3.0 YEAR ONE STATUS

This investigation was to have centered on the development of an IAMS system. The system was to have consisted of computational and informational building blocks, and was to have been tested on flight vehicle models.

Figure 1 shows a schematic of the IAMS system. The model-following intelligent control system has at its core specialized computational and informational units (called building blocks) that may be used in the selection of desirable (if not optimal) aircraft maneuvers. In the final development of the system, informational building blocks will be developed from a priori knowledge. The utility of these building blocks should be demonstrated off-line on a simplified aircraft (A/C) model. The IAMS system, if ultimately completed, will manipulate the building blocks in a way so as to simulate learning in the system. Additionally, the system is designed with the capability of employing an evolutionary algorithm for developing new and highly-effective building blocks on-line. The evolutionary algorithm includes utility measures, performance criteria, constraints, and other important characteristics of the problem environment. Ultimately, the evolutionary algorithm will be used to drive the resulting system's quest for optimal flight maneuvers. Immune features should be incorporated into the system to provide the system with the ability to self-monitor, a mechanism for fault tolerance handling, and a crude memory. Decisions regarding flight maneuvers would ultimately be used to drive an aircraft. Of course, initial testing will require a second aircraft model be used for the purpose of demonstrating the effectiveness of the IAMS system.

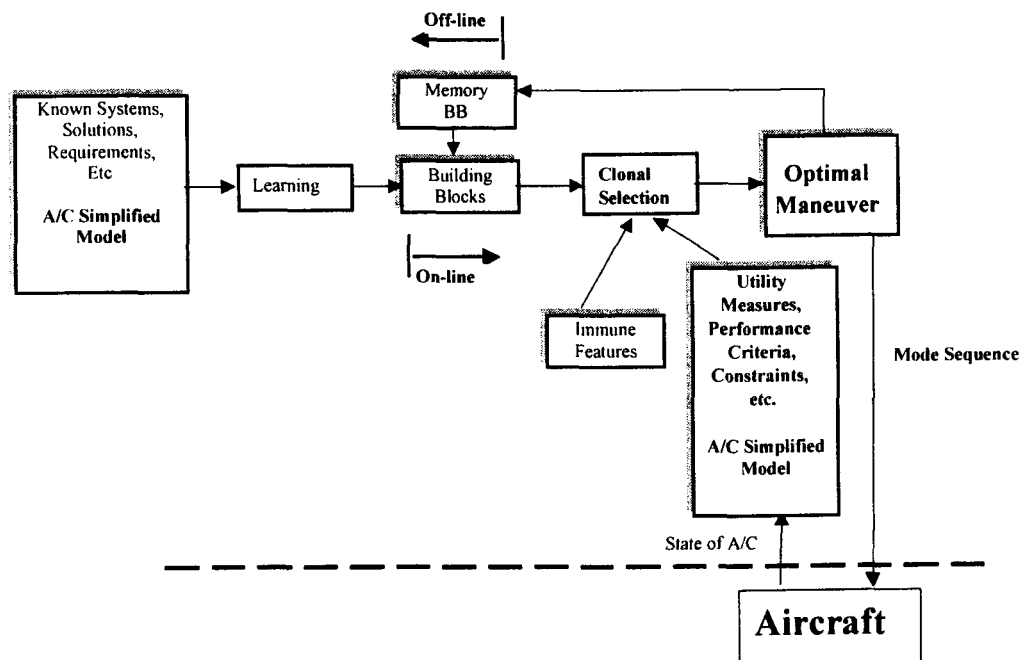


Figure 1: The schematic above shows the overall framework of the IAMS system.

The original proposal for this work described five major tasks to be completed in the development of the IAMS system. These five tasks are:

- (a) determination of the building block architectures;
- (b) identification of immune system features for the maneuver selection process;
- (c) choosing suitable test functions and problems that will be appropriate for solution using an artificial immune system;
- (d) software development for a generic immune system based optimizer;
- (e) integration and testing using vehicle models.

Each of the five tasks is described in more detail later in this section.

It was originally estimated that this project would take two years to complete. The first year effort was to focus on developing an effective building block architecture, immune system characteristic selection, suitable test function selection, and software development. The second year effort was to focus on the particular application of flight maneuver selection. Table 1 shows a schematic of the research schedule as originally proposed. The project has been terminated after the first year.

Table 1: Proposed research schedule.

Task	Months 1-3	Months 4-6	Months 7-9	Months 10-12	Months 13-15	Months 16-18	Months 19-21	Months 22-24
Building Block Architecture Design	X	X						
Immune System Feature Selection		X	X					
Choose suitable test problems		X	X	X				
Software Development			X	X				
IAMS System Testing				X	X	X	X	X
Report							X	X

Efforts did not begin in earnest on this effort until May 2002. Although the grant was awarded in March, the graduate student who is funded by the project was already committed to teaching, grading, and tutoring responsibilities for the Spring 2002 semester within the Aerospace Engineering and Mechanics Department at The University of Alabama. However, the project is very close to being on schedule after the first year despite the slow start. The remainder of this section describes the status of work on the five specific tasks associated with the project at the end of the first year of effort.

3.1 Task 1: Determination of the Building Block Architectures

One of the critical issues associated with the development of the IAMS system ultimately developed will be the identification of effective building blocks. Both informational and computational building blocks will be required. The informational building blocks will consist of known systems, solutions, and requirements for maneuver selection. The computational building blocks will consist of algorithms that can be used in the process of maneuver determination. Further, these building blocks will be manipulated by an evolutionary algorithm and operated on by an artificial immune system. Thus, the architecture (data structure) used to represent the building blocks will need to be powerful and robust. The selection of an appropriate architecture will be critical.

Several critical building blocks were identified for use in the IAMS. These building blocks are segments of a computational system topology that have proven effective in computational modeling and control problems in the past. Examples of building blocks be used include: (1) portions of Fourier series functions and their associated coefficients, (2) neural network connections with their associated weights, and (3) a subsystem of fuzzy logic-based rules.

At the end of year one, portions of Fourier series functions and associated coefficients were incorporated into a comprehensive AIS architecture. The system developed is robust enough to handle the eventual incorporation of more complex building blocks.

3.2 Task 2: Identification of Immune System Features

AISs have developed to the point that they can be quite complex. One important step in the design of an effective IAMS system is to identify the necessary components of an AIS, and to develop computer software that allows for these features to be employed effectively, yet robustly to a wide variety of problems. A basic AIS architecture was developed and tested. The AIS appeared to be robust enough both to be used in the IAMS, and to allow for the incorporation of building blocks necessary to complete the effort.

3.3 Task 3: Choosing Suitable Test Functions

A suite of suitable test functions are to be chosen. Here, the challenge is to identify problem types that are perhaps more appropriately solved using an AIS than using an evolutionary algorithm (EA). Several researchers have studied the characteristics of EA-hard problems, but little has been done to investigate the appropriateness of addressing these problems with an AIS.

One of the main advantages of an AIS over an evolutionary algorithm is an ability both to monitor and to improve a system. The AIS architecture developed under this effort and presented in this report is capable of utilizing receptor-antibody pairs in conjunction with the ability to distinguish self from non-self to effectively monitor and manipulate a system. The effectiveness of the AIS in this domain is later demonstrated via its application to a simple problem of system identification. Although the system identification problem is appropriate for AIS application, other, more general classes of problems must be identified.

3.4 Task 4: Software Development for a Generic Immune System Based Optimizer

Once the building block architecture was designed and the appropriate immune system features were identified, the next step of the project was to incorporate these components into a comprehensive, generic IAMS system. Accomplishing this task will ultimately involve interfacing the artificial immune system software, evolutionary algorithm software, and flight simulation software. At the end of the year one effort, an AIS architecture had been developed, implemented into computer software, and tested on a simple problem. Initial results indicate that the system is flexible enough both to be effectively applied to a flight maneuver selection problem and to allow for the incorporation of more complex informational building blocks. This task was well in hand at the end of year one, but it still remains to be seen as to whether the system will in fact be robust enough to handle the more complex IAMS application.

3.5 Task 5: Integration and Testing Using Vehicle Models

The final step in the IAMS system development was to have been testing. The plan was to test the effectiveness of the system on vehicle models simulations. This was to be the

crux of the second year effort. Naturally, this effort was not completed due to the termination of the project.

3.6 Summary of Status of Research Tasks

When initially proposed, the main efforts of the first year were scheduled to be focused on the development and coding of effective building blocks. These building blocks were to be later incorporated into an AIS, applied to a aircraft maneuver selection problem, thereby forming the IAMS system. As the work was initiated, it was determined that the major focus of the initial efforts needed to be focused more on the development of a generic, robust AIS system. Thus, the main focus of the one year effort was to develop an effective AIS architecture – one that included an evolutionary algorithm and the ability to manage complex building blocks – that was suitable to serve as the heart of the IAMS system. Such an architecture was proposed and tested on a straightforward system identification problem.

Of the four tasks originally defined, three were virtually completed at the end of the first year, and the other two were still in development. First, the critical aspects of artificial immune systems were identified. Second, these attributes were incorporated into a computer software system. Third, the AIS was tested on a simple system identification problem. Of these efforts, really only the expansion of the problem set to which the system was to be applied remained. The focus of the project in the second year was to shift to what was originally thought to be the initial effort: determining effective building blocks (and their respective codings) appropriate for use in the IAMS. Of course if ultimately completed, the IAMS system will also require testing on a flight vehicle simulation.

In summary, despite the fact that efforts on the project were ongoing for only nine months, progress in year one was near the status originally proposed, especially when the fact that the selection of appropriate building blocks was delayed due to logistic considerations.

The remainder of this report is meant to provide a full accounting of efforts in year one. Background information on artificial immune systems is provided, a generic AIS architecture is described, an elementary class of building blocks is presented, and the effectiveness of the system developed to date is demonstrated via its application to a system identification problem.

4.0 BASICS OF THE NATURAL IMMUNE SYSTEM

Organisms in nature rely on their immune systems to defend themselves against unwelcome intruders such as bacteria, potentially dangerous viruses, and other pathogens they encounter on a daily basis. In addition, organisms can develop abnormal cells which can ultimately be lethal if they transform into cancerous cells, and are not removed from the organism. Complex natural organisms have developed effective immune systems that consist of two cooperative defense systems that counter these threats. One of these systems is nonspecific in that it does not distinguish one infectious agent from another.

This nonspecific system includes two lines of defense which a potential invader encounters in sequence. The first line of defense is the external epithelial tissues that cover and line the organism (skin and mucous membranes), and the secretions these tissues produce. The second line of nonspecific defense is internal. This internal second line of defense is triggered by chemical signals and uses antimicrobial proteins and phagocytic cells that indiscriminately attack any invader that penetrates the organism's outer barriers. Inflammation within an organism is an indication that this second line of defense has been deployed.

Although the nonspecific system outlined above can effectively prevent a wide variety of infections, it is the other major defense system, typically referred to as the immune system, which is of more interest in the current context. This third line of defense comes into play simultaneously with the second line of nonspecific defense. Unlike the nonspecific defense, the immune system provides a specific and targeted response to the particular type of invader it faces. This immune response includes the production of specific defensive proteins called antibodies.

Without a doubt, the natural immune system is a complex dynamic system that relies on the interaction of a wide variety of cells derived from the white blood cells (these cells are called lymphocytes). Despite this complexity, it is really the immune system's ability to produce highly specific responses (*specificity*) to a wide variety of invaders (*diversity*), its ability to respond more rapidly to invaders it has incurred previously (*memory*), and its ability to distinguish between cells of the organism and invading molecules (*self/non-self recognition*) that provide its power. In fact, it is exactly these abilities and attributes that Artificial Immune Systems (AISs) seek to replicate. The remainder of this section provides a fairly detailed description of the natural immune system. It is important to again point out the fact that AISs do not represent an attempt to exactly duplicate the operation of the natural immune system, rather they are an attempt to incorporate various characteristics and attributes of the natural immune system into an effective problem-solving paradigm.

4.1 Humoral Immunity and Cell-Mediated Immunity

The natural immune system can actually mount two different types of responses to invaders or disturbances (commonly referred to as antigens): (1) a humoral response and (2) a cell-mediated response. Humoral immunity results in the production of antibodies, which are secreted by certain lymphocytes and circulate as soluble proteins in blood plasma and lymph, fluids that were long ago called humors. Around the turn of the century, researchers transferred such fluids from one animal to another and found that this process transferred immunity. However, the researchers also found that immunity to some conditions could be passed along only if lymphocytes were transferred. This second type of immunity, which depends on the direct action of cells (certain types of lymphocytes) rather than antibodies, became known as cell-mediated immunity.

Thus, the immune system is made up of two branches. The circulating antibodies of the humoral branch defend mainly against toxins, free bacteria, and viruses present in body

fluids. In contrast, lymphocytes of the cell-mediated branch are active against bacteria and viruses inside the host's cells and against fungi, protozoa, and worms. Cell-mediated immunity is also involved in attacks on transplanted tissue and cancer cells, both of which are perceived as foreign substances or invaders, e.g., they are considered to be "non-self" as opposed to "self." Thus, both transplanted tissue and cancer cells can be attacked by the cell-mediated branch of the immune system.

4.2 Cells of the Immune System

Organisms ranging from the most developed to the very simple are known to have effective immune systems. Here the focus is on the most complex immune system, that of mammals. The immune systems of most mammals are known to be populated by two main classes of lymphocytes: B cells (B lymphocytes) and T cells (T lymphocytes). The B cells carry out the humoral immune response while the T cells are used mainly in the cell-mediated immune response. Like all blood cells, lymphocytes originate from pluripotent stem cells either in the bone marrow or in the fetus (mainly in the liver). Initially, all lymphocytes are the same. They later differentiate into T cells or B cells depending on where they ultimately mature. Lymphocytes that migrate from the bone marrow to the thymus (a gland in the upper region of the chest) develop into T cells. B cells, on the other hand, remain in the bone marrow where they complete their maturation.

Once the B and T cells mature, they are typically stored in the lymph nodes, the spleen, and other lymphatic organs. It is through the system of lymph nodes that these B and T cells are most likely to encounter antigens. Both B and T cells are equipped with very specific antigen receptors on their plasma membranes. Recalling that the B cells are generally used in a non-specific response, their antigen receptors are actually membrane-bound antibody molecules specific for a certain antigen. The antigen receptors of T cells are different from antibodies, but these T cell receptors recognize antigens as specifically as antibodies do. Thus, the specificity and diversity of the immune system both depend on receptors on each B cell and T cell that enable that lymphocyte to identify and respond to a particular antigen.

When the receptor of a lymphocyte binds to an antigen, the lymphocyte is activated to divide and differentiate. The result is a population of effector cells and it is these cells that actually defend the body in an immune system response. In a humoral response, B cells activated by the binding of an antigen gives rise to effector cells called plasma cells. These plasma cells secrete antibodies that help eliminate the particular antigen to which the B cells are bound. The operation of a cell-mediated response is different. There are actually two different types of effectors in a cell-mediated response: (1) cytotoxic T cells (T_C) and (2) helper T cells (T_H). The cytotoxic T cells kill infected cells and cancer cells. The helper T cells, on the other hand, secrete protein factors called cytokines. Cytokines are molecules that are secreted by one cell as a regulator of neighboring cells. Cytokines actually help regulate B cells and T cells, and therefore play a pivotal role in both humoral and cell-mediated responses.

4.3 Clonal Selection is Basis for Specificity and Diversity

The ability of the natural immune system to respond to the millions of potential antigens it could face is a remarkable aspect of the system. In fact, this *diversity* is one of the key factors that AISs will need to replicate. It is easy to envision at least two mechanisms by which a natural immune system can accomplish this task. In the first scenario, each lymphocyte could be thought of as plastic in its response – it has the ability to tailor its action to match whatever antigen it is battling at a given time (changing the antibody it secretes depending on the invader with which it is faced). In the second scenario, a lymphocyte might be designed to respond to a single antigen with one particular antibody. The first scenario would arguably require lymphocytes with some sort of “intelligence” while the second scenario would require the immune system to have a very large collection of lymphocytes. Evidence supports the second scenario: natural immune systems depend on an enormous diversity of antigen-specific lymphocytes.

Each lymphocyte’s *specificity* for an antigenic target is rigidly predetermined during embryonic development, before any encounter with a given antigen ever takes place. The mark of this specificity is the antigen receptor the lymphocyte bears on its surface. The lymphocyte may or may not ever come into contact with the corresponding antigen. If that antigen does enter the body and binds to receptors on the specific lymphocytes, then those (and only those) lymphocytes are activated to mount an immune response. Basically, a collection of highly specific-lymphocytes are stored in the organism.

When a lymphocyte is activated, it proliferates by cell division and develops into a large number of identical effector cells – clones of the lymphocytes are produced that combat the very antigen that provoked the response. For example, plasma cells that develop from an activated B cell all secrete the same type of antibody that functioned as the antigen receptor on the original B cell that first encountered the antigen. This antigen-specific selection and cloning of lymphocytes is called *clonal selection*. The concept of clonal selection is so fundamental to understanding immunity that it is worth restating: each antigen, by binding to specific receptors, selectively activates a tiny fraction of cells from the body’s diverse pool of lymphocytes. This relatively small number of selected cells gives rise to a clone of millions of effector cells, all dedicated to eliminating the specific antigen that stimulated the humoral or cell-mediated immune response.

4.4 Memory Cells for Secondary Immune Responses

The primary immune response is based on the selective proliferation of lymphocytes to form clones of effector cells upon first exposure to an antigen. This initial response to an antigen typically takes approximately 5 to 10 days for maximum production of effector cells. During this lag period, the lymphocytes selected by the antigen are differentiating into effector T cells and antibody-producing plasma cells. However, the immune system has the ability to improve its performance because of a kind of memory intrinsic to the system. For example, if an organism is exposed to the same antigen at some later time, the response is much faster and more prolonged than the primary response. This

secondary response results in maximum production of effector cells on subsequent exposures is on the order of 3 to 5 days. Aside from an improved response time, the secondary response is also more effective than the primary response in that the antibodies produced are more effective in binding to the antigen than those produced during the primary response.

The immune system's ability to recognize an antigen as previously encountered is called *immunological memory*. This ability is based on long-lived memory cells, which are produced along with the relatively short-lived effector cells of the primary immune response. During the primary response, these memory cells are not active. However, they survive for long periods and proliferate rapidly when exposed again to the same antigen that caused their formation. The secondary immune response gives rise to a new clone of memory cells, as well as to new effector cells. It is by this mechanism that childhood exposure to diseases such as chickenpox usually confers immunity for a lifetime.

4.5 Molecular Markers Provide Self/Non-self Recognition

Aside from *specificity* and *diversity*, another key feature of the natural immune system is its ability to distinguish between "self" and "non-self." This feature is critical in natural immune systems because without it the immune system may and does actually attack the organism itself. The antigen receptors on the surfaces of lymphocytes are responsible for detecting foreign molecules that enter the body. Normally, there are no lymphocytes that are reactive against the body's own molecules. Self-tolerance begins to develop as T and B lymphocytes bearing antigen receptors mature in the thymus and bone marrow, and continues to develop even as the cells migrate to lymphoid tissues. Any lymphocytes with receptors for molecules present in the body are destroyed or are rendered nonfunctional leaving only lymphocytes that are reactive against foreign molecules.

Among the important "self markers," the native molecules tolerated by an individual's immune system, are a collection of molecules encoded by a family of genes called the major histocompatibility complex (MHC). In humans, the MHC is also called the HLA (human leukocyte antigen) group. These genes encode glycoproteins (proteins with carbohydrate attached) embedded in the plasma membranes of cells. Because there are at least 20 MHC genes and as many as 100 alleles for each gene, it is virtually impossible for any two people, except identical twins, to have matching sets of MHC markers on their cells. Thus, the major histocompatibility complex is a biochemical fingerprint unique to each individual.

Two main classes of MHC molecules mark cells as "self." Class I MHC molecules are located on all nucleated cells – almost every cell of the body. Class II MHC molecules are restricted to a few specialized cell types of the body's defense system, including macrophages, B cells, and activated T cells. Class II MHC molecules play an important role in interactions between cells of the immune system.

4.6 B Cells Generate Specific Antibodies

Although most AISs are typically based on the cell-mediated response, the humoral response is described in more detail for completeness. Recall that humoral immunity results in the production of antibodies, which are secreted by certain lymphocytes and circulate as soluble proteins in blood plasma and lymph, fluids that were long ago called humors. In practice this actually occurs when B cells with specific receptors (membrane-bound antibodies) are stimulated by an antigen to differentiate into a clone of plasma cells, which begin to secrete antibodies. Antibodies are most effective against pathogens that are circulating in the blood and lymph, such as free bacteria and viruses. This selective activation of B cells arms the body with long-lived memory cells that function in a secondary humoral response. What distinguishes this branch of the immune system is its production of antibodies, which are not involved in cell-mediated defense. Notice, however, that the two branches (humoral and cell-mediated) often operate together, so that the distinctions between the two branches sometimes break down.

4.7 The Activation of B Cells

In most cases, the selective activation of a B cell to form a clone of plasma cells and memory cells is a two-step process. One step, discussed earlier, is the binding of antigen to specific receptors on the surface of the B cell. The other step involves two other types of cells: (1) macrophages and (2) helper T cells. After a macrophage engulfs pathogens by phagocytosis, fragments of the partially digested antigen molecules are bound by class II MHC molecules. The two molecules are transported to the plasma membrane and displayed on the surface of the macrophage, which is then known as an antigen-presenting cell (APC). A helper T cell's specific receptors recognize this self/non-self combination of MHC and a particular antigen fragment. Antigen-specific contact between the T cell and the antigen-presenting macrophage activates the T cell, which proliferates and forms a clone of helper T cells keyed to the specific antigen. The helper T cells then secrete cytokines, which selectively stimulate B cells that have already encountered that particular antigen. A helper T cell contacts a B cell in the same way it contacts a macrophage displaying the antigen. When antigen receptor on a B cell binds to an antigen the cell takes in a few of these foreign molecules by endocytosis. The B cell then displays antigen fragments bound to class II MHC markers on the cell surface. The helper T cell's receptor recognizes and binds to this antigen-MHC complex.

Thus, macrophages and B cells both act as APCs in their interactions with helper T cells, but there is an important difference: each macrophage can display a wide variety of antigens, from the many different pathogens it has taken in by phagocytosis. But, each B cell, being specific, can bind to and subsequently display only the antigens of a particular pathogen. Helper T cells are also antigen-specific and can be activated only by those macrophages presenting the appropriate antigen in combination with a class II MHC molecule. In this way, the nonspecific defense afforded by macrophages also enhances the specific defense by selectively priming T cells. The activated T cells then stimulate the appropriate B cells to mount a humoral immune response against a certain antigen.

4.8 T-Dependent and T-Independent Antigens

Antigens that evoke the response described in Section 2.7 are known as T-dependent antigens because they cannot stimulate antibody production without the involvement of T cells. Most antigens are T-dependent. However, certain types of antigens, called T-independent antigens, trigger humoral immune responses without involving macrophages or T cells. These antigenic molecules are usually long chains of repeating units, such as polysaccharides or proteins with many similar polypeptide subunits. Such antigens are found in bacterial capsules and bacterial flagella. Apparently, the numerous subunits of such antigens bind simultaneously to a number of the antigen receptors on the B cell surface. This provides enough of a stimulus to the B cells without the assistance of T cells. However, the humoral response (antibody production) to T-independent antigens is generally much weaker than the response to T-dependent antigens. Furthermore, no memory cells are generated in T-independent responses.

Once activated by T-dependent or T-independent antigens, a B cell gives rise to a clone of plasma cells, and each of these effector cells secretes as many as 2000 antibodies per second for the 4- to 5-day lifetime of the cells. These specific antibodies help eliminate the foreign invader. The response to T-dependent antigens also generates a clone of long-lived memory cells.

4.9 The Molecular Basis of Antigen-Antibody Specificity

Most antigens are proteins or large polysaccharides. These molecules are often outer components of the coats of viruses, the capsules and cell walls of bacteria, and the surface molecules of many other types of cells. Foreign molecules associated with transplanted tissues and organs or with blood cells from other individuals or species can also incite an immune response. The surfaces of foreign substances such as pollen also include antigens.

Antibodies do not generally recognize the antigen as a whole molecule. Rather, they identify a localized region on the surface of an antigen called an antigenic determinant, or epitope. A single antigen such as a bacterial protein may have several effective epitopes, stimulating several different B cells to make distinct antibodies against it. Different parts of the bacterial cell may possess different antigens. Thus, the immune system responds to a particular species of bacterial cell by producing many antibodies specific for the various epitopes marking the antigens of the bacterial cell wall, capsule, and flagella.






Antibodies constitute a class of proteins called immunoglobulins (Igs). Every antibody molecule has at least two identical sites that bind to the epitope that provoked its production. A typical antibody has four polypeptide chains joined to form a Y-shaped molecule: two identical light chains and two identical heavy chains. Both heavy and light chains have constant regions. These regions are called constant because their sequences of amino acids vary little among the antibodies that perform a particular type of defense, despite a wide variation in antigen specificity. At the tips of the Y-shaped molecule's two arms are the variable regions of the heavy and light chains, so named because their amino

acid sequences vary extensively from antibody to antibody. These tips of the Y function as the antigen-binding sites. The dimensions and contours of the binding sites are determined by the unique amino acid sequences of the variable regions of the heavy and light chains. The association between an antigen-binding site and an epitope resembles that between an enzyme and its substrate: several weak bonds form between contiguous chemical groups on the respective molecules.

The antigen-binding site is responsible for an antibody's recognition function – its ability to identify a specific epitope of an antigen. The tail of the antibody molecule, consisting of the constant regions of the polypeptide chains, is responsible for the antibody's effector function, the mechanism by which it inactivates or helps destroy an antigenic invader.

There are five types of constant regions, and these determine the five major classes of mammalian immunoglobulins: IgM, IgG, IgA, IgD, and IgE. The shapes and functions of these basic antibody types are summarized in Table 2. Each class is characterized by a type of constant region that enables the antibody molecules to perform certain defense functions. For example, IgA can be transported across epithelia, and it is present in saliva, sweat, and tears. Within each Ig class is an enormous variety of specific antibodies with unique antigen-binding sites.

Table 2: Five Basic Antibody Types – There are five basic types of antibodies: IgM, IgG, IgA, IgD, and IgE. Their basic properties are summarized in this table.

The Five Classes of Immunoglobulins	
 <p>IgM (pentamer)</p>	<p>IgMs are the first circulating antibodies to appear in response to an initial exposure to an antigen. Their concentration in the blood declines rapidly. This is diagnostically useful because the presence of IgM usually indicates a current infection by the pathogen causing its formation. IgM consists of five Y-shaped monomers arranged in a pentamer structure. The numerous antigen-binding sites make it very effective in agglutinating antigens and in reactions involving complement. IgM is too large to cross the placenta and does not confer maternal immunity.</p>
 <p>IgG (monomer)</p>	<p>IGG is the most abundant of the circulating antibodies. It readily crossed the walls of blood vessels and enters tissue fluids. IgG also crosses the placenta and confers passive immunity from the mother to the fetus. IgG protects against bacteria, viruses, and toxins circulating in the blood and lymph, and triggers action of the complement system.</p>
 <p>IgA (dimer)</p>	<p>IgA is produced primarily in the form of two Y-shaped monomers (a dimer) by cells abundant in mucous membranes. The main function of IgA is to prevent the attachment of viruses and bacteria to epithelial surfaces. IgA is also found in many body secretions, such as saliva, perspiration, and tears. Its presence in colostrums (the first mile of a nursing mammal) helps protect the infant from gastro-intestinal infections.</p>
 <p>IgD (monomer)</p>	<p>IgD antibodies do not activate the complement system and cannot cross the placenta. They are mostly found on the surfaces of B cells, probably functioning as an antigen receptor required for initiating the differentiation of B cells into plasma cells and memory B cells.</p>
 <p>IgE (monomer)</p>	<p>IgE antibodies are slightly larger than IgG molecules and represent only a very small fraction of the total antibodies in the blood. The tail regions attach to receptors on mast cells and basophils and, when triggered by an antigen, cause the cells to release histamine and other chemicals that cause an allergic reaction.</p>

4.10 How Antibodies Work

An antibody does not usually destroy an antigenic invader directly. The binding of antibodies to antigens to form an antigen-antibody complex is the basis of several effector mechanisms. The simplest of these is neutralization, in which the antibody blocks certain sites on an antigen, making it ineffective. Antibodies neutralize a virus by attaching to the sites the virus must use to bind to its host cell. Similarly, coating a

bacterial toxin with antibodies effectively neutralizes it. Eventually, phagocytic cells dispose of the antigen-antibody complex.

Another effector mechanism is the agglutination (clumping) of bacteria by antibodies. Agglutination is possible because each antibody molecule has at least two antigen-binding sites and can cross-link adjacent antigens. The bacterial clumps are easier for phagocytic cells to engulf than are single bacteria. A similar mechanism is precipitation, the cross-linking of soluble antigen molecules (rather than cells) to form immobile precipitates that are captured by phagocytes.

One of the most important effector mechanisms of the humoral response is the activation of the complement system by antigen-antibody complexes. Recall that the complement system is a group of proteins that acts cooperatively with elements of the nonspecific and specific defense systems. Antibodies often combine with complement proteins, activating the complement proteins to produce lesions in a foreign cell's membrane that result in lysis (bursting) of the cell.

4.11 The Activation of T Cells

In contrast to B cells, T cells cannot detect free antigens present in body fluids. T cells respond only to antigenic epitopes displayed on the surfaces of the body's own cells. These bound antigens are recognized by specific T-cell receptors embedded in the plasma membranes of T cells. One example is the binding of helper T cells to antigens displayed on the surfaces of antigen-presenting macrophages and B cells in the activation of humoral immunity. The receptor of a T_H cell actually recognizes a combination of an antigen fragment and one of the body's self-markers, class II MHC, on the surface of a macrophage or B cell. In contrast, the receptor of a T_C cell recognizes a combination of an antigen fragment and a class I MHC molecule, a self-marker found on every nucleated cell of the body. In both cases, the MHC protein that presents the antigen is shaped something like a hammock that nestles the antigen. An MHC molecule can associate with a variety of antigens. However, each combination of MHC and antigen forms a unique complex that is recognized by specific T cells.

The interaction between a T_H cell and an antigen-presenting cell (APC) is greatly enhanced by the presence of a T-cell surface molecule called CD4. Present on most helper T cells, CD4 has an affinity for a part of the class II MHC molecule. The CD4-class II MHC interaction helps keep the T_H and the APC engaged while antigen-specific contact is going on. Similarly, a cytotoxic T cell bears surface molecules called CD8, which help in the interaction with class I MHC molecules.

The MHC-antigen complex is like a red flag to T cells, calling them into action against cells infected by the pathogen represented by that particular antigen. The situation stimulates T cells that have the appropriate receptors to proliferate and form clones to activated T_H or T_C cells specialized for fighting that particular pathogen.

One function of helper T cells is their role in activating B cells to secrete antibodies against T-dependent antigens during humoral responses. T_H cells also activate other types of T cells to mount cell-mediated responses to antigens.

The ability of helper T cells to stimulate other lymphocytes depends on cytokines. As a macrophage engulfs and presents antigen, it is stimulated to secrete a cytokine called interleukin-1. This signals the T_H cell to release another cytokine known as interleukin-2. In an example of positive feedback, interleukin-2 stimulates T_H cells to grow and divide more rapidly, increasing the supply of both T_H cells and interleukin-2. The interleukin-2 and other cytokines released by this growing population of T_H cells also help activate B cells, thus stimulating the humoral response against a specific antigen. And cytokines from T_H cells help arm the cell-mediated response by stimulating T_C cells.

4.12 How Cytotoxic T Cells Work

Cytotoxic T cells kill host cells infected by viruses or other intracellular pathogens. Such infected cells display antigens complexed with class I MHC molecules. Notice an important difference in the functions of the receptors on T_H and T_C cells: the receptors on helper T cells enable those lymphocytes to stimulate immunity by binding specifically to antigen-presenting cells marked by class II MHC molecules, such as macrophages and B cells. In contrast, cytotoxic T cells, by recognizing specific antigens in association with class I MHC markers, can bind to any cell of the body infected by that particular antigenic invader. When a T_C cell binds to an infected cell, it releases perforin, a protein that forms an open lesion in the infected cell's membrane. The target cell loses cytoplasm through the lesion, leading to cell lysis. While this destroys the host cell, it also deprives the pathogen of a place to reproduce and exposes the pathogen to circulating antibodies. This cell-mediated action is an essential feature of the body's overall defense system, because antibodies cannot attack pathogens that have already invaded host cells. After destroying an infected cell, the T_C cell continues to live and can kill many other cells.

Cytotoxic T cells also function in defense against cancer. Cancer cells arise periodically in the body and, because they carry distinctive molecular markers not found on normal cells, are identified as non-self by the immune system. T_C cells target the cancer cells and lyse them. Cancers are more likely to become established in individuals with defective immune systems and elderly people with declining immunity.

Although T_C cells and natural killer cells attack target cells by similar mechanisms, there is an important difference between the roles of these cells in defense: natural killer cells, part of the body's nonspecific defense, do not respond to specific antigens.

A third type of T lymphocyte, called a suppressor T cell (T_S) is not well understood. (Some immunologists believe that T_S cells are actually a type of T_H cell rather than a separate subclass of T cells.) T_S cells probably function in turning off the immune response when an antigen is no longer present (hence their name).

4.13 The Immune System's Capacity to Distinguish Self from Non-Self

In addition to distinguishing between an animal's own cells and pathogens such as bacteria and viruses, the immune system wages war against cells from other individuals of the same species. For example, a skin graft from one person placed on another person (not an identical twin) will look healthy for a day or two, but after that it will be destroyed by the immune system. It is interesting that a mother does not reject a fetus as a foreign body. Apparently, the structure of the placenta, the link between mother and fetus, is the key to this acceptance; the mother rejects fetal tissue placed elsewhere in her body.

There is a large body of information about the potential problems associated with blood transfusions and organ transplants, two important applications of the immune system's response to non-self.

5.0 ARTIFICIAL IMMUNE SYSTEM

The natural immune system is extremely powerful, and very effective at monitoring, adapting, and defending the organism in which it resides and operates. As demonstrated in the previous section, the natural immune system consists of numerous entities and subsystems. When developing an artificial immune system, it is important to recognize that the goal is not to exactly duplicate the operation of the natural immune system, rather it is to borrow ideas from the natural counterpart to produce a computational tool for solving, in this case, engineering problems.

When developing a computational model of a natural system, it is reasonable to consider the minimum components from the natural counterpart required in the computational system. Thus, this section considers the key features of the natural immune system. Later, computational models of these features are presented.

5.1 Key Features of the Natural Immune System

The natural immune system is capable of developing specific, effective responses against a wide range of disturbances including foreign microbes, toxins, and even transplanted tissue. Four key features characterize the natural immune system: (1) specificity, (2) diversity, (3) memory, and (4) self/non-self recognition.

5.1.1 Specificity

The natural immune system is capable of recognizing and combating particular microorganisms and foreign molecules. A foreign entity that elicits this type of response from the natural immune system is termed an antigen. The natural immune system responds to specific antigens by activating lymphocytes and by producing special proteins called antibodies. Some of the antigens that regularly trigger an immune system response include viruses, bacteria, fungi, protozoa, and parasitic worms. Of course, there are other things that can trigger an immune system response such as pollen, insect venom,

and transplanted tissue. Each antigen has specific attributes (typically shape) and triggers the production of the exact type of antibody necessary to combat it. Thus, each response of the natural immune system is targeted toward a particular disturbance. This ability to distinguish between various invaders, and the ability to secrete the appropriate antibody, is termed specificity. Certainly, this is a key feature of the natural immune system.

5.1.2 Diversity

The natural immune system is capable of responding to literally millions of kinds of disturbances, each identified by its specific antigenic markers. This ability to respond appropriately to such a wide array of invaders is possible because the immune system relies on an enormous store of lymphocyte populations; each population bears receptors for a particular antigen. Each population of lymphocyte is stimulated by a specific antigen and responds by producing and administering the proper type of antibody. This capability is known as diversity.

5.1.3 Memory

The natural immune system has the ability to “remember” antigens it has encountered previously, and to react to them more promptly and more effectively on subsequent encounters. This feature, commonly called acquired immunity, was recognized as early as 400 BC when Thucydides of Athens described how those sick and dying during an epidemic of plague were cared for by those who had recovered, “for no one was ever attacked a second time.” This concept of immunity is familiar to most. For example, if a person had chicken pox as a child, they will not have the disease a second time. This acquired immunity can be thought of as a form of memory.

5.1.4 Self/Non-Self Recognition

One of the more important features of the natural immune system is its ability to distinguish between “self” and “non-self.” Without this capability, the natural immune system would attack itself. In fact, organisms in which this capability is not present are thought of as having autoimmune disorders, and can have their own tissues destroyed by their immune system. Obviously, this is a critical feature of a natural immune system.

5.2 Key Features of an Artificial Immune System

There have been a number of architectures proposed for AISs. Despite the varying nature of these architectures, each seems to incorporate (or at least acknowledge the potential for) features that allow for the AIS: (1) to be capable of identifying and combating specific disturbances or problems (where the disturbances and/or problems are thought of as antigens), (2) to have the means to effectively address a wide variety of disturbances, (3) to possess some capacity for memory, and (4) to be able to distinguish between a normal or acceptable mode of operation (self) and a “disturbed state” (non-self). These capabilities have been instantiated in systems that have been given numerous

names. However, they can typically be described by the following sub-systems (Dasgupta, 1999): (1) bone marrow models, (2) negative selection models, (3) clonal selection algorithms, (4) immune network models, and (5) immunized computational systems. Systems proposed and implemented in the past have proven to be effective in the domain to which they have been applied. However, none of them have been shown to be robust enough to incorporate the complexity required of the current IAMS system, e.g., the ability to incorporate computational or information al building blocks. The AIS architecture described in the next section of this report has the potential to be such a system.

5.3 Architecture for an Artificial Immune System

This section provides an overview of the AIS architecture proposed for the current effort; one that incorporates the capabilities cited in Section 5.2 above. This architecture is an extension of the architecture originally proposed by Forrest et al (1994) and implemented on a tool breakage problem by Dasgupta and Forrest (1996). This proposed algorithm relies on receptors to identify various classes of problems to be solved within the confines of a system, antibodies consisting of various combinations of building blocks to address problems attacking the system, an evolutionary algorithm to explore potentially effective antibodies, a database of antibodies demonstrated to be effective against known antigens, and a mechanism for recognizing self from non-self.

At its highest level the architecture is capable of distinguishing between the normal mode of operation of a system (self) and a disturbed mode of operation (non-self). It includes a library of receptors that are used to monitor the system. Appended to these receptors are combinations of building blocks that form the equivalent of antibodies which serve as solutions to various problems that can occur within the system. When a problem (antigen) is identified (a non-self state), the coded non-self state is compared to the library of known receptor-antibody combinations, and the one that is most closely related to the current problem is cloned. The population of clones is adapted using an evolutionary algorithm until a satisfactory antibody is developed. At this point, the system is altered, and the newly developed effective antibody is placed in the library of effective receptor-antibody combinations.

In this approach, the receptors are analogous to the lymphocytes of the natural immune system. These are used to distinguish between self and non-self which amounts to a crude negative selection algorithm. The coded strings appended to the receptors are analogous to antibodies, and the building blocks that compose the antibodies are akin to the proteins that make up natural antibodies. The library of effective solutions effectively provides the AIS with a form of memory. The solutions contained in the library could be allowed to exist on a schedule similar to that of a natural immune system, but this capability has not proven to be necessary when used in the problems of the scope of those presented later in this paper. The algorithm for this AIS is summarized as:

- (a) Define self as a collection S of strings of length L over a finite alphabet, p .
- (b) Generate a set of R receptors, each of which fails to match any string in S . The receptors are over an alphabet $p+1$, where the additional character in the alphabet is a wild-card character that matches any symbol in S . These receptors are each of length k , yet have appended to them a set of building blocks which are coded representations of pieces of potential solutions to the problem or antigen.
- (c) Monitor S for changes by continually matching the receptors in R against S . If any receptor is matched, then a change is known to have occurred.
- (d) If a receptor is matched exactly, then go to the library of known effective solutions and use this solution to perturb the system back to S . Else, clone the matched receptors to create a population of antibodies. Employ an evolutionary algorithm to generate an effective antibody for the current antigen or state of the system. Once an effective antibody is determined, add the receptor-antibody combination to the existing library.
- (e) Continue with step (c) above.

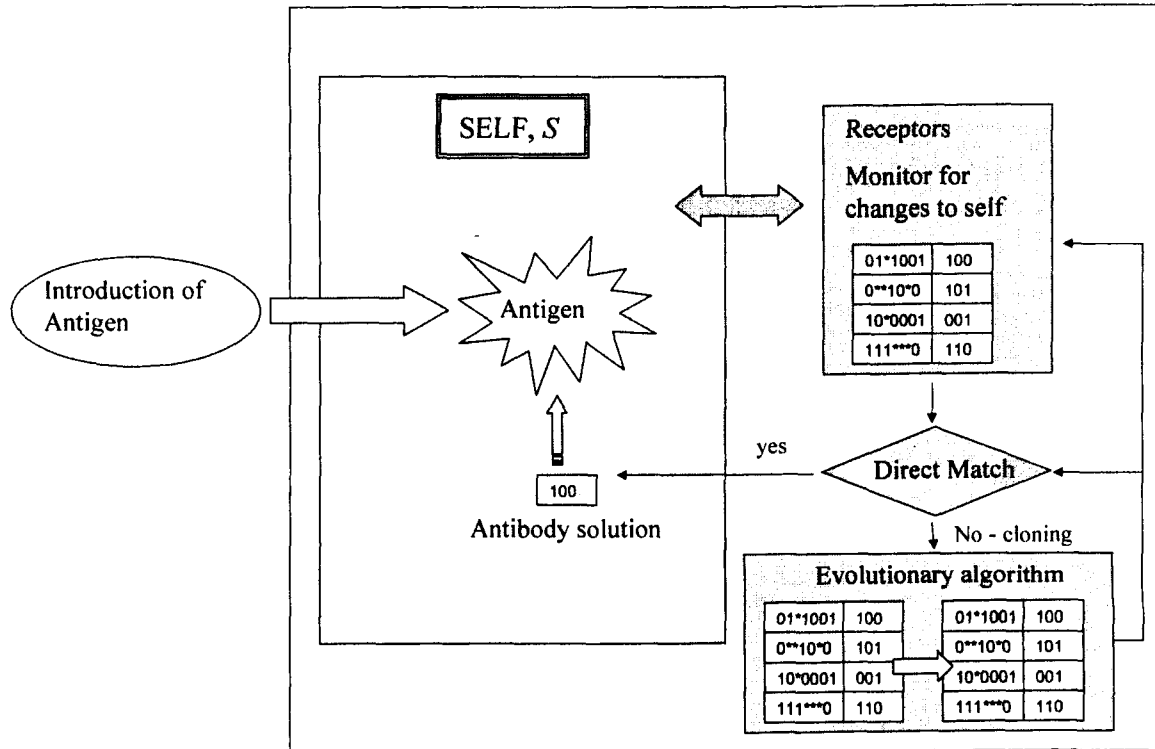


Figure 2: The AIS algorithm is depicted below. Once an antigen is introduced, receptors, antibodies, and an evolutionary algorithm are used to repair the system.

Figure 2 depicts the AIS architecture graphically. "Self" is typically some system being operated in "the world." It may be any of a number of things including a control system manipulating a robot, a network of computers running in a constrained environment, or a system for storing medical records. During the course of normal operation of this system,

there is some mechanism by which a disruption can occur; perhaps it is a natural degradation of one of the components of the system, or maybe a change prescribed by an external user of the system. The system is continually monitored by a set of receptors. The negative selection algorithm responsible for this monitoring process compares the state of the system known as "self" to a set of receptors designed to identify "non-self" states of the system. Once an antigen is known to exist, the non-self state is compared to a library of receptors, each of which has an associated antibody or fix to the problem incurred by the system. If a receptor is matched exactly, then the known solution to the problem is sent to the system which then returns to "self." However, if no known solution is know (no receptor from the library matches exactly), then an evolutionary algorithm is used to design an effective solution by combining building blocks into a comprehensive solution. When the evolutionary algorithm completes its task, the newly-located effective antibody is sent to the system which returns to "self," and the receptor-antibody pair are added to the library of know effective solutions.

Like with virtually all AISs, this architecture depends on a priori knowledge of the system on which the AIS is being applied. Details of two specific implementations of this algorithm are presented in the remainder of this document. The two problems discussed are both related to adaptive control systems: (1) a system identification problem and (2) an adaptive control system.

6.0 SYSTEM IDENTIFICATION PROBLEM

Over the last decade many researchers have focused on the development of adaptive, intelligent control systems. The expert systems of the artificial intelligence movement of the late 1960's and 1970's have given way to techniques from the field of computational intelligence (neural networks, genetic algorithms, learning classifier systems, fuzzy logic, etc.) of the 1990's, and many successful applications of intelligent control systems have been presented (KrishnaKumar, Kaneshige, and Satyadas, 2002). These adaptive systems rely on a variety of computational techniques to provide computer control systems with the characteristics of human controllers including abstraction, adaptability, and reasoning-like capabilities. The robust control systems are beginning to find their way into mainstream engineering applications.

Many of the most successful intelligent control systems depend on some sort of computational model of the system being controlled (the plant). This dependency has necessitated an increased effort in solving what is known as the system identification problem: developing a computational model of the plant, at times with incomplete knowledge of all of the plant, i.e., one does not always know the complete equations of motion or operation of the system being controlled. Basically, this problem boils down to the ability to "predict" the response of a system given various inputs. Many adaptive control systems use this type of plant model to investigate new and novel control strategies, or to tune existing control laws in real-time. This section describes the application of an AIS to a simple system identification problem.

6.1 Overview of the System Identification Problem

The system identification problem presented here is sometimes known as “function-matching.” The idea is that a generic function is being monitored, and time-series values are recorded for use within the system. The objective is to utilize this time-series data to accurately “predict” future responses of the system. Figure 3 shows an example of time-series data that can be used for a system identification problem. Here, the objective is to develop a model of the function $f(t)$ such that future values can be accurately predicted. This does not appear to be terribly difficult problem here because the oscillatory nature of the data makes it likely that the function is of the form $f(t) = c_1 \sin(c_2 t) + c_3 \cos(c_4 t)$. In fact, the function used is $f(t) = 1.0 \sin\left(\frac{\pi}{2} t\right) + 1.0 \cos\left(\frac{\pi}{2} t\right)$, and there is a variety of methods that can be used to effectively solve this system identification problem.

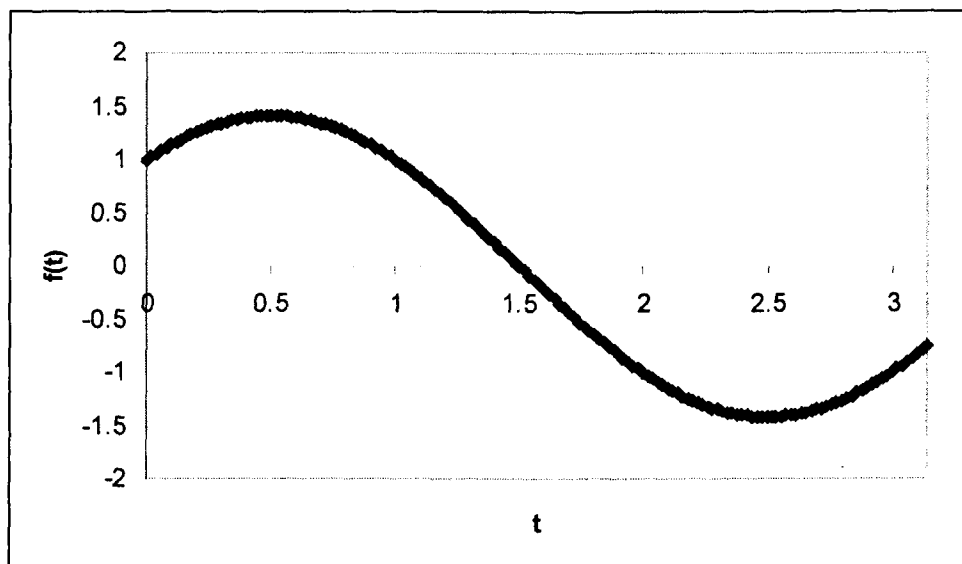


Figure 3: The system identification problem requires that some sort of system model be created so that the response of the system can be accurately predicted.

Figure 3 demonstrates one aspect of a system identification problem. However, the problem can become markedly more difficult when the system dynamics change with time. Figure 4 shows a sample function in which the model changes over three distinct regions: (1) $0 \leq t < 3.14$, (2) $3.14 \leq t < 6.28$, and (3) $6.28 \leq t < 10.0$. Here, the system identification problem requires that the function be matched for any given time. Now, an automated system capable of modeling the function requires at least two capabilities: (1) monitoring the function to determine when it has changed, and (2) updating the system model so that the response of the system can be predicted. This seems to be an ideal application of an AIS.

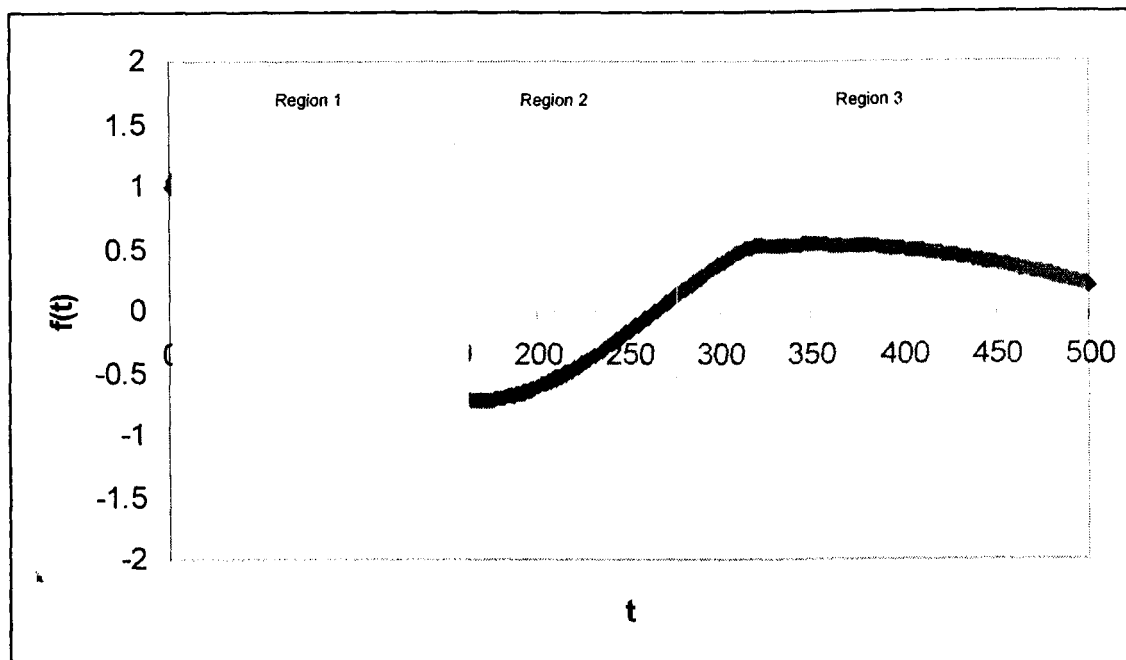


Figure 4: The ability to accurately predict the response of the function above requires both monitoring to determine when the system has changed, and updating to produce a refined model of the system.

6.2 Self vs. Non-self in the System Identification Problem

The ability of the AIS to distinguish between self and non-self is the responsibility of the negative selection algorithm. In the context of the system identification problem, this ability is tied to the idea of monitoring the function for changes in the system characteristics. Bluntly, the negative selection algorithm must provide a mechanism by which the AIS can determine whether or not a new modeling equation is required.

A key concept in the ability to distinguish between self and non-self is to equate self to “normal operation” of the system being monitored (the system to which the AIS is being applied). For instance, assuming that the time-series response of a system shown in Figure 3 is accurately modeled at time $t=0$, then the negative selection algorithm of the AIS must be able to recognize that the system has changed at time $t=150$, and it must be able to do so with no other information than a sequence of time-series values.

Here, the focus will be on providing a scheme by which the normal response of the system can be represented in such a way as to interact with the receptors, antibodies, and evolutionary algorithm to be described in the following sections. The idea is to encapsulate the performance of the system into a collection of binary strings that can be considered representative of the normal operation of the function. This will be accomplished in three distinct steps. First the range of the system response (say $-1.5 \leq f(t) \leq 1.5$) for the function shown in Figure 4 is divided into a discrete number of “bins.” Of course, the greater the number of bins, the more accurately the function can be

represented. Once the space has been discretized, then each point can be thought of as falling into an individual bin. Each bin will be identified with a binary number, so that it has a unique tag. Figure 5 provides a visual depiction of this operation as applied to a sin function. Here, note that there are $2^4 = 16$ distinct bins – an integer power of 2, represented as p), is used so that each bin will have a unique binary identifier that is p boolean characters long.

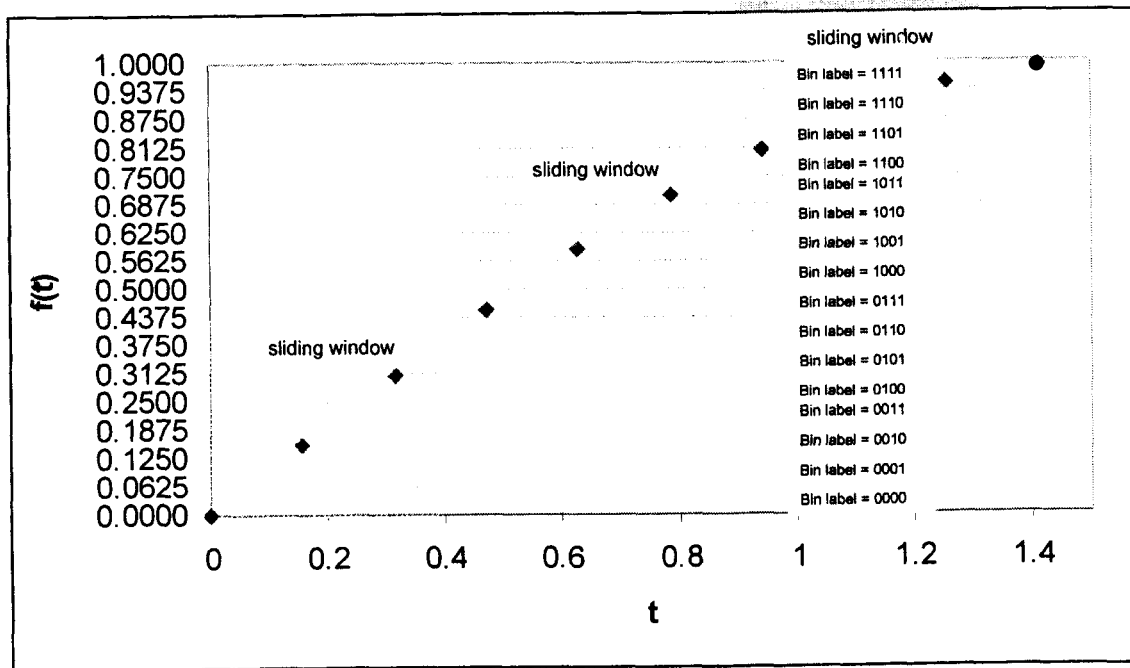


Figure 5: The range of the function is divided into a discrete number of “bins,” each of which has a unique identifier.

The second step involves grouping data values in the time series. After the space of the function has been divided into a discrete number of bins, a sequence of data values, say $k=3$ for instance, can be grouped and represented by k bits of length 2^p . For example, the first three data values from Figure 5 would be identified as: 000000100100 (the first point falls in bin “0000”, the second falls in bin “0010”, and the third falls in bin “0100”).

The final step in the definition of self by the negative selection algorithm is to prepare a database of the strings described above. When database is complete, the result can be thought of as a table of binary strings such as that shown in Table 3 below. It is important to note that the accuracy of this definition of self is improved as (1) the discretization of the space if made finer, (2) the number of points gathered in a window is “appropriate,” and (3) the time over which the data is gathered increases. Note also that the sliding window can be moved a total of k time steps, or it can be moved incrementally one data value at a time.

Table 3: This table is indicative of the definition of self generated by the negative selection algorithm designed for the AIS.

Function Value 1	Function Value 2	Function Value 3	Identifying String
0.0000	0.1564	0.3090	000000100111
0.4540	0.5878	0.7071	101111011110
0.8090	0.8910	0.9511	111111111111
0.9877	1.0000	0.9877	111111111111
0.9511	0.8910	0.8090	111111101100
0.7071	0.5878	0.4540	101110010111
0.3090	0.1564	0.0000	010000100000

6.3 Receptors for Monitoring System Performance

Once a database of binary strings has been created that represents self, there are at least two approaches to determining when the system upon which the AIS is operating is no longer in a normal mode of operation. First, the current k -sized data set is gathered from the system. This data set is then converted into a string, and the string is compared to the strings in the definition of self. If the string representing the current data set matches a string in self, then the system is in a normal mode of operation. This direct approach is manageable, but does not fit with the idea of receptors-antibody pairs of the AIS.

The second approach to identifying normal modes of operation relies on a database of receptors. These receptors, like the strings identifying self, are of length $p*k$. These receptors are strings that represent data sequences that are not indicative of a normal mode of operation. The current k -sized data set can again be coded into a string representation, and compared to the library of receptors. Since it would be difficult to produce an all-inclusive library of receptors that represent all possible combinations of non-self, a wild-card character (*) is introduced. The wild-card is interpreted to match either a "0" or a "1" in any string. Thus, the receptor 00000011111* would match both the string 000000111110 and the string 000000111111. In this way, a smaller library of receptors can be used to identify non-self modes.

The receptors must be generated such that no string in the library of receptors matches any string in the database of strings defining self. The receptors are generated randomly from the alphabet (0,1, *). After N receptors are generated, they are placed in a library for comparison to the current status of the system. Each sequence of data values from the system (here, a sequence is p data values so that the corresponding coded strings will be of the same length as both strings identifying self and the receptors) is then compared to the strings in the receptor library. If a data sequence matches a receptor, then it can be inferred that the system is no longer operating in a normal mode. This indirect mechanism of comparing the operation of the system to a library of non-self receptors is more in line with the operation of a natural immune system, and thus is employed in the AIS being described.

6.4 Receptor-Antibody Pairs

The receptors are used to monitor the status of the system: when data sequences from the system match a receptor, then it is inferred that the system is no longer in a normal mode of operation, e.g., an antigen is present. This process allows the AIS to achieve one very important aspect of its operation – identifying when an antigen is present. The next step, of course, is to provide a mechanism by which the AIS can eliminate the antigen. It is the receptor-antibody pairs that make this possible.

Each receptor has attached to it an antibody. The antibody is a portion of a bit-string that represents a potential solution to the problem posed by the antibody. Each antibody is represented by a sub-string from the alphabet $V\{0,1\}$. In the system identification problem, it is assumed that the function to be matched can be represented by a Fourier series. That is, it is assumed that $f(t) = \sum_{i=1}^N A_i \sin(\omega_i t) + B_i \cos(\omega_i t)$, where A_i , B_i , and ω_i are coefficients that must be determined such that $f(t)$ accurately depicts the performance of the function (or system). Thus, each sub-string representing an antibody must encode values for all coefficients A_i , B_i , and ω_i necessary to depict the function. When a receptor is matched by a data sequence from the system, its attached antibody represents values of the coefficients in the Fourier series necessary to accurately depict the new state of the system.

As an example, consider the simplest case in which it is known that the function $f(t)$ is described by the equation $f(t) = A \sin(\omega t) + B \cos(\omega t)$. Here, each antibody must contain a bit-string representation of the terms A_i , B_i , and ω_i . If we choose to represent each coefficient with say four bits, then a receptor-antibody pair might appear as the string shown in Figure 6.

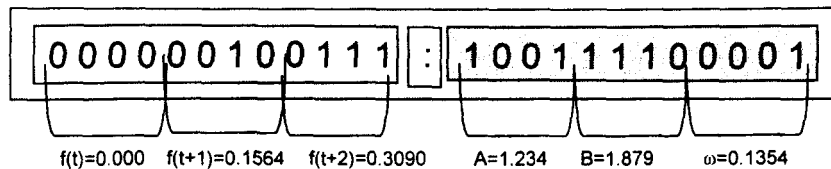


Figure 6: A receptor-antibody pair will represent both a data sequence from the system and coefficients that can be used to construct a new and accurate model of the system.

Note in Figure 6 that the receptor depicted is an exact match with a condition in the system (no “*” characters appear in this example). In this case, the implication is that this is a scenario for which a known solution exists. It is in this way that memory is added to the system: as known effective solutions to specific system conditions are generated by the evolutionary algorithm, they remain in the library of receptors and are used to immediately update the system model.

With the introduction of receptor-antibody pairs, the AIS now has the ability to: (1) monitor the status of the system by comparing data sequences from the system to a library of receptors and (2) correct the system model for known data sequences by using antibodies attached to receptors for known conditions. What is left to do is to establish some mechanism by which the AIS can explore potentially effective solutions to situations that are not known a priori: a mechanism for locating effective solutions to unseen problems is required. The mechanism proposed here is a form of an evolutionary algorithm.

6.5 Evolutionary Algorithm for Performing Clonal Selection

The receptors in the AIS are composed of both "0" and "1" characters, but also of the wild card character "*." In this way, generality can be incorporated into the system: it is easier to produce a set of receptors that match a wider variety of non-self conditions than if only the alphabet $V\{0,1\}$ is used. Unfortunately, the fact that each receptor includes an antibody from the alphabet $V\{0,1\}$ means that a single antibody could be described for a variety of system states. Naturally, a specific antibody may not be suitable for a wide range of antigens. Thus, some mechanism is required to locate effective antibodies for any antigen that could occur. The exploratory algorithm employed in the present system is an evolutionary algorithm.

Evolutionary algorithms are search routines based on the mechanics of natural genetics. They combine a Darwinian survival-of-the-fittest approach with a structured information exchange. In their rudimentary form, they manipulate bit-strings. Their ability to locate near-optimal solutions quickly makes them an ideal choice for the task at hand.

The process invoked here is to:

- (1) identify all receptors that are matched by the current state of the system;
- (2) generate an initial population for manipulation by the evolutionary algorithm by mutating each of the antibodies of the matched receptors until the requisite number of strings is created;
- (3) invoke a run of the evolutionary algorithm until a suitable antibody is determined (in this case, until a set of coefficients are found that provide for an accurate model of the function);
- (4) determine the exact state of the system that matched a receptor;
- (5) create the string that is an exact receptor;
- (6) add this string to the definition of self and start re-generating a current definition of self;
- (7) store the receptor-antibody pair in the receptor library;
- (8) wait for the next data sequence from the system.

This process is analogous to the clonal selection algorithm apparent in natural artificial immune systems. Although it is not an exact duplicate, the efficient nature of evolutionary algorithms should be appropriate for use in completing the task of generating an effective solution (antibody) to a given problem (antigen). The following

section describes the results of using this AIS architecture in the system identification problem.

6.6 Results of Using the AIS on a System Identification Problem

The AIS was used to predict the response of a function $f(t) = a \sin(bt) + c \cos(dt)$. This function was altered every 50 time units in accordance with the parameter values shown in Table 4. Thus, the AIS was faced with the dual task of identifying when the function had changed (every 50 time units), and employing receptor-antibody pairs in conjunction with the evolutionary algorithm to identify new values of the parameters. Of course, the objective is to locate new parameter values that allow for accurate replication of the function. Figure 7 shows the function over the 200 time-unit span.

Table 4: Parameter values used to define the target function $f(t)$

Time	a	b	c	d
$0 \leq t < 50$	1.0	$\frac{\pi}{2}$	1.0	$\frac{\pi}{2}$
$50 \leq t < 100$	1.0	π	1.0	π
$100 \leq t < 150$	2.0	π	2.0	π
$150 \leq t < 200$	2.0	2π	2.0	2π

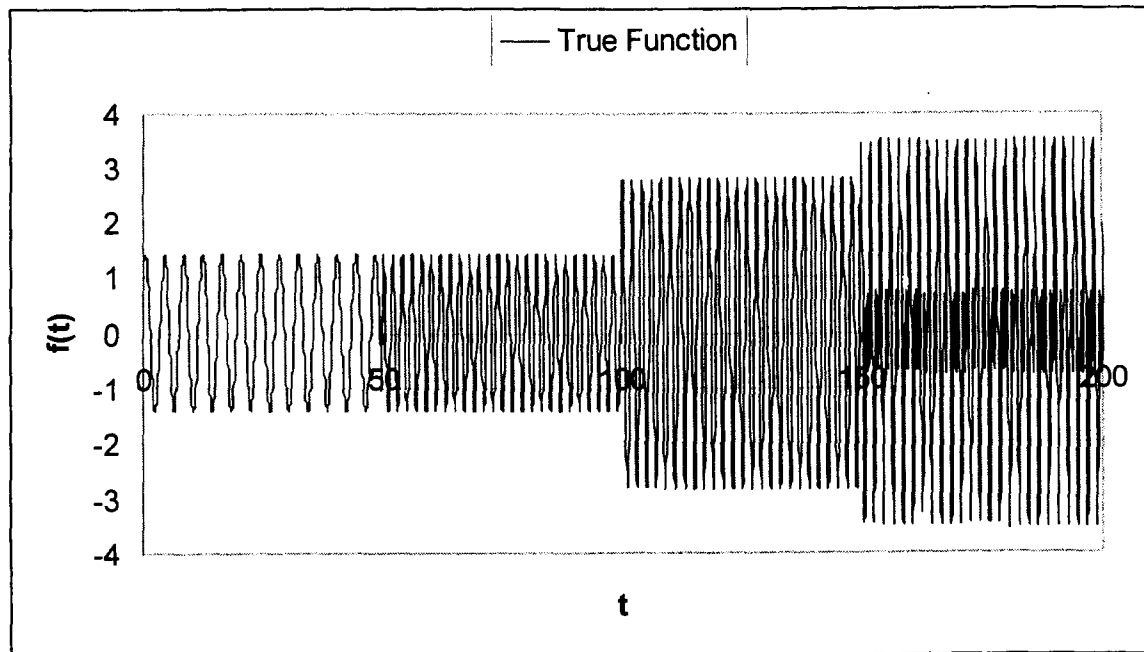


Figure 7: The function changes every 50 seconds as shown above.

The first requirement of the AIS is to provide identify when a change to the function has occurred. This task is accomplished through the matching or “firing” of receptors that identify “non-self” conditions. Figure 8 shows the number of receptors fired at given times. Note that as the simulation progresses, the number of receptors fired grows from three to eight. Thus, the AIS is fully capable of identifying that a change has occurred; a stimulus the initiates the search for accurate function parameter values using an evolutionary algorithm.

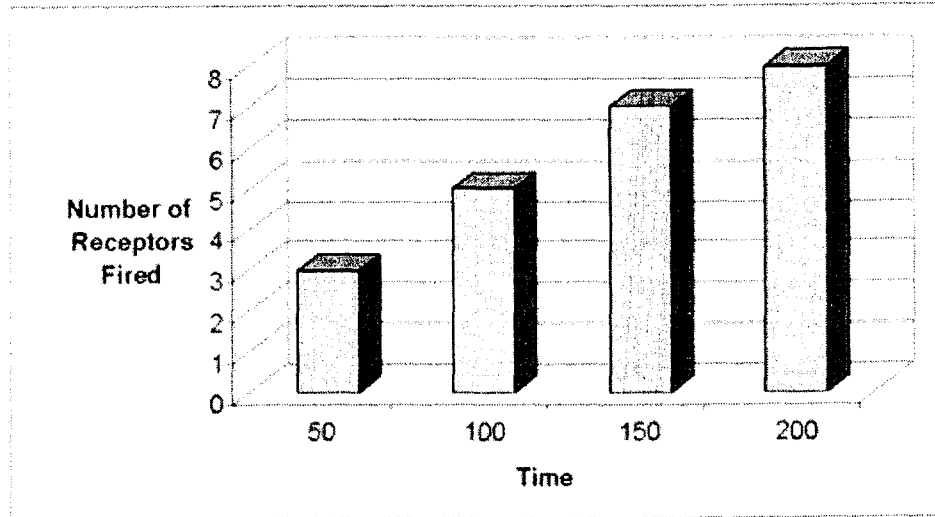


Figure 8: The number of receptors fired during the simulation grows from three to eight over the course of the simulation.

Figure 9 shows the performance of the AIS in matching the function over the first 55-second period. Note that the function has obviously changed at the 50-second mark, and that the AIS is no longer predicting the function values well. Figure 10 shows that it takes the AIS’s evolutionary algorithm approximately 10 seconds to identify valid parameters for the function. Once these parameters are located (an antibody is found), the AIS predicted function once again accurately matches the true function. It is important to note that no attempts were made to optimize the operation of the evolutionary algorithm with respect to time, e.g., the process was run in real-time in conjunction with the function definition. Naturally, this time delay can be reduced with the implementation of parallel processing.

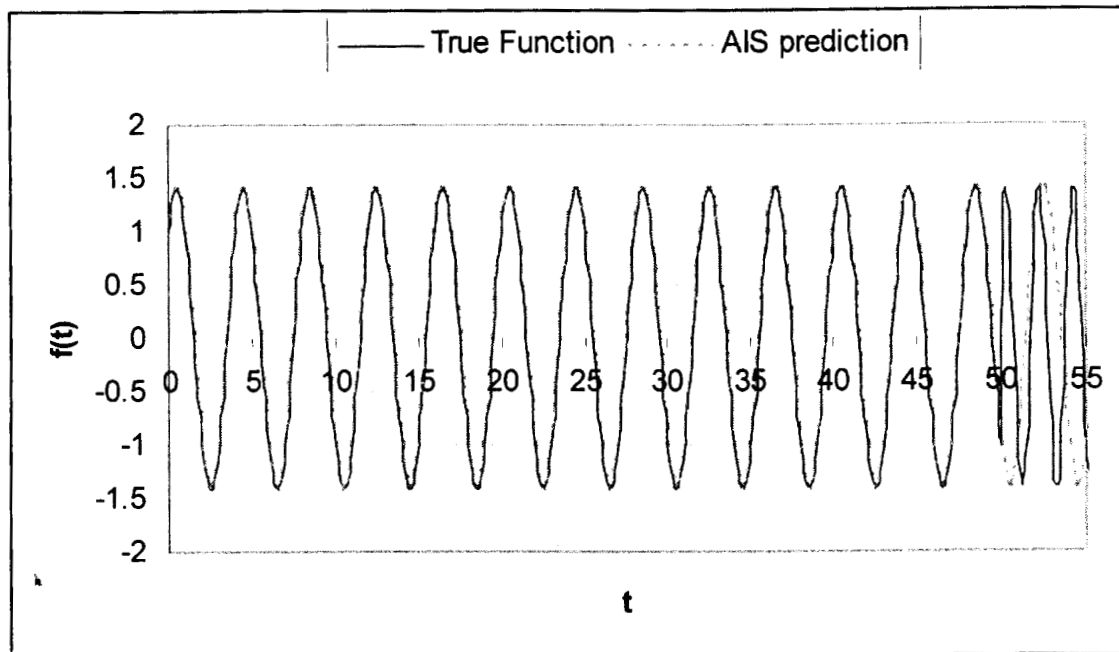


Figure 9: The function is altered at $t=50$ seconds, and the AIS-predicted response is no longer accurate.

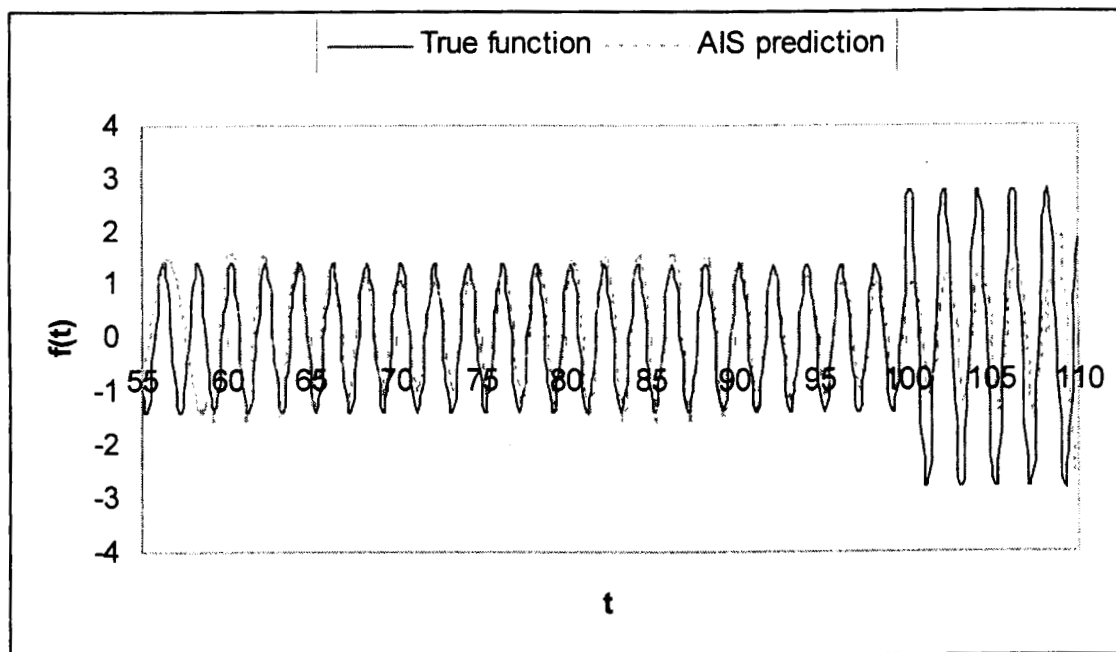


Figure 10: After a 10 second period in which the AIS is searching for an accurate solution, the predicted function once again matches the true function.

Figures 11 and 12 show the operation of the AIS over the second and third 50-second time intervals. As with the initial invocation of the search algorithm, the AIS is able to identify that the function has been changed, and to locate parameters (antibodies) that produce an accurate response.

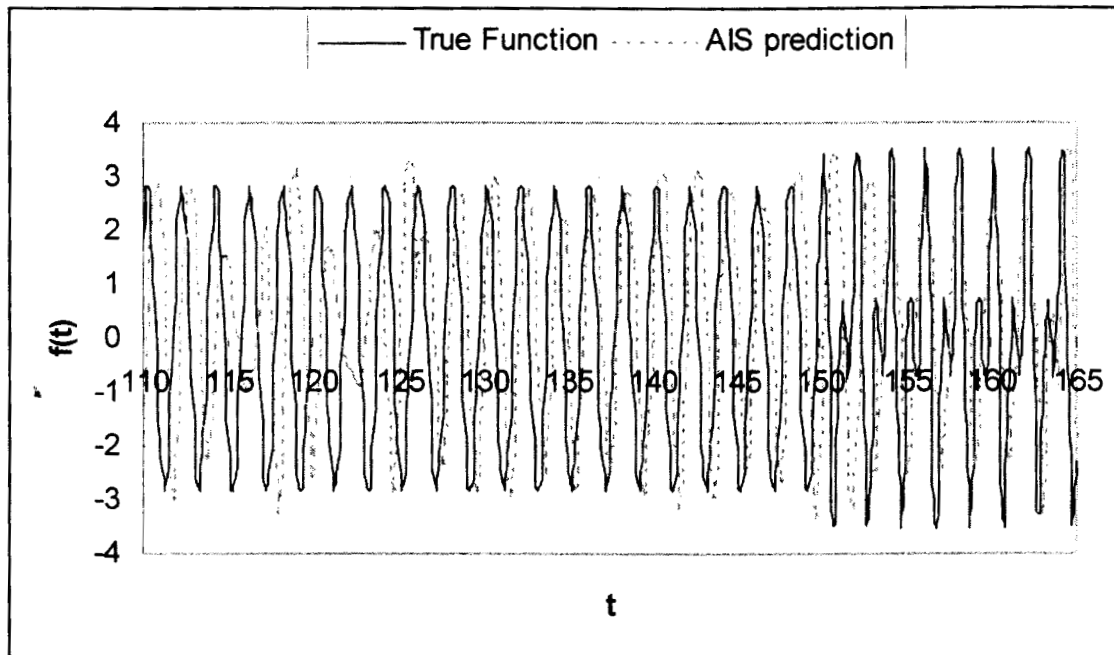


Figure 11: AIS solution over the second 50-second period of operation.

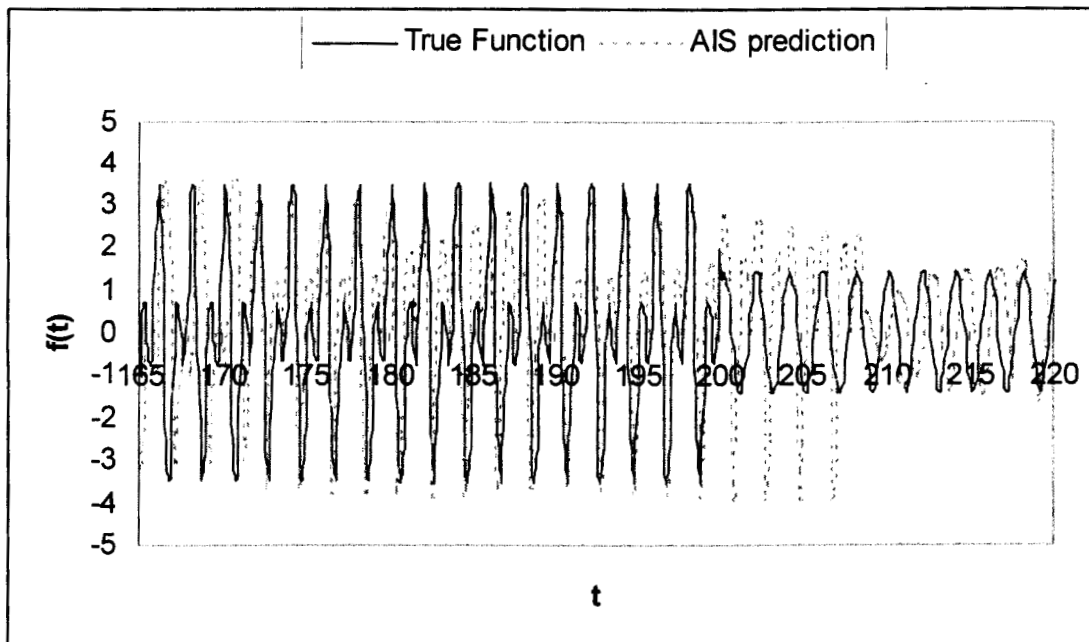


Figure 12: AIS solution over the third 50-second period of operation.

This section has provided a brief glimpse of the effectiveness of the proposed AIS architecture's performance in a simple function matching problem. Results indicate that this architecture is capable of monitoring and matching the response of the system. This simple system identification example will be extended in future efforts. In addition, the AIS architecture will be applied in an environment in which the objective is not function identification, rather the objective will be to adapt a control system manipulating a dynamic system.

7.0 STEPS FOR COMPLETING THE PROJECT

Two key elements of the five key tasks originally proposed are yet to be completed. First, the most appropriate building block structures (and associated codings) suitable for the IAMS should be determined. Although it is proposed that sub-structures of neural networks, fuzzy logic-based rules, and Fourier functions and associated coefficients will be adequate to effectively solve the flight maneuver selection problem, this supposition is yet to be demonstrated. Once the appropriate building blocks are identified, they could easily be incorporated into the existing AIS architecture. Of course, the final step would be to test the entire system on vehicle simulation software currently being used at NASA-Ames Research Center.

Ultimate completion of the project, if pursued should focus on the two areas stated above. It is expected that the selection and incorporation of suitable building blocks would be approximately a six-month effort. The testing of the IAMS would likely require a full six months also. These efforts could reasonably be expected to be completed with the funding not exceeding that originally proposed.

8.0 REFERENCES

- Bäck, T., Fogel, D. B., & Michalewicz, Z. (Eds.) (1997). *Handbook of evolutionary computation*. New York: Oxford University Press.
- Dasgupta, D. (Ed.) (1999). *Artificial immune systems and their application*. New York: Springer-Verlag.
- Dasgupta, D., & Forrest, S. (1996). Tool breakage detection in milling operations using a negative-selection algorithm. University of New Mexico Department of Computer Science Report. Technical Report No. CS95-5. Albuquerque, NM, August 15, 1995.
- Fiesler, E., & Beale, R. (Eds.) (1997). *Handbook of neural computation*. New York: Oxford University Press. 02-212.
- Forrest, S., Perelson, A.S., Allen, L., & Cherukuri, R. (1994). Self-nonsel self discrimination in a computer. IEEE Symposium on Research in Security and Privacy, IEEE Computer Society Press, Los Almitos, CA, pp.

Hofmeyr, S. A. (2000). An interpretative introduction to the immune system. In I. Cohen & L. A. Segel (Eds.), *Design principles for the immune system and other distributed autonomous systems*, New York: Oxford University Press.

Hofmeyr, S. A., & Forrest, S. (1999). Immunity by design: An artificial immune system. Proceedings of the Genetic and Evolutionary Computing Conference'99, 1289-1296, Orlando, FL.

Jungheui, C., & Deaton, R. (2001). Introduction to DNA computing. 2001 Genetic and Evolutionary Computation Conference Tutorial Program, 480-492, San Francisco, CA.

KrishnaKumar, K., Kaneshige, J., & Satyadas, A. (2002). Challenging Aerospace Problems For Intelligent Systems, Presented at the von Karman Lecture series on Intelligent Systems for Aeronautics, Belgium, May 2002.